Distributed Quantum Computing

Harry Buhrman and Hein Röhrig*

Centrum voor Wiskunde en Informatica (CWI)
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
Harry.Buhrman@cwi.nl, hein@acm.org

Abstract. Quantum computing combines the framework of *quantum mechanics* with that of computer science. In this paper we give a short introduction to quantum computing and survey the results in the area of distributed quantum computing and its applications to physics.

1 Introduction

Computing is a physical process and therefore the theory of computing should incorporate the laws of physics. Quantum mechanics, developed during the last century, is to date the most accurate description of nature. Quantum computing is the area that combines the laws of quantum mechanics with computer science. In this paper we give a short introduction to quantum computing and its formalism; for a more detailed treatment of this we refer the reader to the excellent textbook of Nielsen and Chuang [46]. Quantum bits or "qubits" are the basic building blocks for quantum computers. As was shown already in the seventies by Holevo [30], qubits *cannot* be used to compress messages better than with bits. In general, a k-bit message needs also k qubits to be stored or sent over a channel. Qubits can, however, reduce the communication of certain distributed computational tasks, as was first demonstrated in [22] and subsequent papers, among them [17,23,21,48,7].

We survey some of these results here. The first result of a cheaper quantum than classical communication protocol [22] was inspired by nonlocality experiments constructed by physicist in order to test the strange and nonlocal behavior of entanglement. In 1935, Einstein, Podolsky, and Rosen devised a thought experiment that sought to show how quantum mechanics is incomplete because it would allow for some form of faster-than-light communication. Much later, in 1964, Bell [10] came up with an experimental way of testing the nonlocal behavior of quantum mechanics. These tests and the so-called Bell inequalities lead to experiments [9], that seem to demonstrate the nonlocality of quantum mechanics. However, these tests suffered from the drawback that implementations in the lab are error prone and sometimes do not give the right outcome or none at all. When the classical local theory is also allowed to make such errors, it can be shown that the nonlocality tests can also be explained by classical

* Supported in part by the EU fifth framework project RESQ, IST-2001-37559, and NWO grant 612.055.001.

B. Rovan and P. Vojtáš (Eds.): MFCS 2003, LNCS 2747, pp. 1–20, 2003.

physics, and do not demonstrate the nonlocal behavior of quantum mechanics at all! In this paper we survey how the results obtained in quantum communication complexity can be used to propose nonlocality experiments that are robust against errors and would for the first time demonstrate conclusively nonlocality. The paper is organized as follows. In Sect. 2 we give a short introduction to quantum mechanics and the notation used in this paper. In Sect. 3 we describe the quantum analogue of the black-box model of computation and describe one of the first quantum algorithms, due to Deutsch and Jozsa. Section 4 surveys some of the results in distributed computing.

2 Quantum Mechanics and Computing

One of the main and very counterintuitive features of quantum mechanics is the *superposition* principle. A physical system may be in a superposition of two or more *different* states at the same time. Quantum mechanics prescribes that when we observe such a system we see one of these states with a certain probability resulting in a collapse of the system into the state that we observed.

2.1 Qubits, Superposition, and Measurement

Let us concentrate now to computation. Classically a bit can be in any of two states: 0 or 1. Quantum mechanically a quantum bit or qubit may be in a superposition of both 0 and 1. It is useful to describe such systems as vectors in a finite-dimensional Hilbert space, in this case a two-dimensional one. We will identify the vector $\begin{bmatrix}1\\0\end{bmatrix}$ with the symbol $|0\rangle$ to denote the classical bit 0 and vector $\begin{bmatrix}0\\1\end{bmatrix}$ with the symbol $|1\rangle$ denoting the classical bit 1. This notation is called Dirac or "ket" notation, from "bra-ket." The "bra" is $\langle\ |$ and $\langle a|b\rangle$ denotes the inner product between a and b. Quantum mechanics now allows for a superposition of these two classical states:

$$\alpha|0\rangle + \beta|1\rangle = \begin{bmatrix}\alpha\\\beta\end{bmatrix} \ , \tag{1}$$

where α and β, called *amplitudes*, are complex numbers with the property that

$$|\alpha|^2 + |\beta|^2 = 1 \ . \tag{2}$$

Next, *observing* or *measuring* a qubit $\alpha|0\rangle + \beta|1\rangle$ will yield outcome 0 with probability $|\alpha|^2$ and 1 with probability $|\beta|^2$. Moreover, after this measurement the qubit is either in the classical state $|0\rangle$ if we measured a 0, and in $|1\rangle$ if we measured a 1. Note that equation (2) guarantees that a qubit, when measured, indeed induces a probability distribution over 0 and 1.

Let us try to plug in some values for α and β:

$$\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle \tag{3}$$

Observing this qubit will result with probability 0.5 in seeing a 0 and with probability 0.5 in a 1. In general, our system will consist of more than just one qubit. Equations (1) and (2) generalize in the obvious way. Suppose we want to model k qubits. Classically k *bits* can be in any of 2^k different configurations: $1 \dots 2^k$. This means that k *qubits* can be in a superposition of all (or part) of these 2^k basis states:

$$\alpha_1 \overbrace{|00\dots0\rangle}^{k} + \dots + \alpha_{2^k} \overbrace{|11\dots1\rangle}^{k} = \sum_{i \in \{0,1\}^k} \alpha_i |i\rangle \tag{4}$$

with the additional requirement that

$$\sum_{i \in \{0,1\}^k} |\alpha_i|^2 = 1 \ . \tag{5}$$

When observing these k qubits we will see i with probability $|\alpha_i|^2$.

If we have two qubits $|x\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$ and $|y\rangle = \beta_0|0\rangle + \beta_1|1\rangle$ then $|x\rangle \otimes |y\rangle$ are the two qubits in a four-dimensional Hilbert space. This construction is called the tensor or Kronecker product:

$$\begin{aligned} |x\rangle \otimes |y\rangle &= (\alpha_0|0\rangle + \alpha_1|1\rangle) \otimes (\beta_0|0\rangle + \beta_1|1\rangle) \\ &= \alpha_0\beta_0|00\rangle + \alpha_0\beta_1|01\rangle + \alpha_1\beta_0|10\rangle + \alpha_1\beta_1|11\rangle. \end{aligned}$$

By convention, $|0\rangle \otimes |0\rangle$, $|0\rangle|0\rangle$, and $|00\rangle$ denote the same thing.

In general not all the two-qubit states that satisfy (4) and (5) are obtained as the tensor of two qubits. We will see an important example, the EPR pair, in Subsect. 2.3. Such states are called *entangled*.

2.2 Unitary Operations

Next we would like to model operations on qubits. Quantum mechanics tells us that these operation have to be modeled as *linear* operations with the additional constraint that these operations preserve the probability interpretation, i.e., the squares of the amplitudes sum up to 1 (see (2) and (5)). Such transformations are called *unitary*; they are the square matrices U that satisfy

$$UU^* = I \ ,$$

where U^* is the complex conjugate transpose of U and I is the identity matrix. In terms of computation, the unitarity constraint implies that the computation is *reversible*.

The following transformation on a single qubit is important and very useful. It is called the Hadamard transform.

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{6}$$

It is a unitary operation since:

$$\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}\cdot\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}=\begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix}$$

Now let us do a Hadamard operation on a qubit that is in the classical state $|0\rangle$:

$$\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}\cdot\begin{bmatrix}1\\ 0\end{bmatrix}=\frac{1}{\sqrt{2}}\begin{bmatrix}1\\ 1\end{bmatrix}$$

This is in ket notation $\frac{1}{\sqrt{2}}|0\rangle+\frac{1}{\sqrt{2}}|1\rangle$, which is the random qubit from (3). When we apply the Hadamard transform again on this qubit,

$$\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}\cdot\frac{1}{\sqrt{2}}\begin{bmatrix}1\\ 1\end{bmatrix}=\begin{bmatrix}\frac{1}{2}+\frac{1}{2}\\ \frac{1}{2}-\frac{1}{2}\end{bmatrix}=\begin{bmatrix}1\\ 0\end{bmatrix}\quad, \tag{7}$$

we get the $|0\rangle$ again. The important point is the minus sign in the Hadamard transform. Its effect is illustrated in (7) above. The minus sign caused the $\frac{1}{2}-\frac{1}{2}$ in the lower half of the vector to cancel out, or interfere destructively, while both terms in the upper half interfered constructively. It is the superposition principle together with this interference behavior that gives quantum computing its power.

The tensor product is also defined on linear operations. If we have an $m\times n$ matrix A and an $m'\times n'$ matrix B then $A\otimes B$ is a $(m\cdot m')\times(n\cdot n')$ matrix defined as:

$$\begin{bmatrix}a_{1,1}\cdot B & a_{1,2}\cdot B & \dots & a_{1,n}\cdot B\\ a_{2,1}\cdot B & a_{2,2}\cdot B & \dots & a_{2,n}\cdot B\\ \vdots & \vdots & \ddots & \vdots\\ a_{m,1}\cdot B & a_{m,2}\cdot B & \dots & a_{m,n}\cdot B\end{bmatrix}$$

2.3 Einstein-Podolsky-Rosen Paradox

In Sect. 2.1 we have seen that any set of k qubits is admissible if it satisfies (4) and (5). Bearing this in mind let us examine the following state consisting out of 2 qubits:

$$\frac{1}{\sqrt{2}}|00\rangle+\frac{1}{\sqrt{2}}|11\rangle \tag{8}$$

Note that the first 0 and the first 1 form the first qubit and the second 0 and the second 1 form the second qubit. This state is called the "EPR state" after its inventors Einstein, Podolsky, and Rosen [25]. The purpose of this state was to devise a thought experiment to show the incompleteness of quantum mechanics. Imagine that we have this EPR state and that Alice has the first qubit somewhere on Mars and that Bob has the second, say, here on earth. If Alice measures her qubit she will see a 0 or a 1 with equal probability and the state will have collapsed to either $|00\rangle$, if she saw a 0 or $|11\rangle$ in case it was a 1. The same is true for Bob. This leads to the following situation. Suppose that the first qubit, on Mars, was measured first and that Alice saw a 1. This now means that when Bob measures his qubit he will also measure a 1. It appears that some information, i.e., the outcome of Alice's measurement, has somehow traveled to

earth *instantaneously*. Since nothing can travel faster than the speed of light something must be wrong.

It turns out that EPR pairs cannot be used for communication: straightforward arithmetic shows that the probabilities of Bob obtaining a certain measurement outcome are not changed no matter what Alice does. However, they can be used to reduce *communication complexity* as we are going to see in Sect. 4.2.

Classical bits can be copied. Qubits on the other hand can *not* be copied.

Theorem 1. *[24,52] Qubits cannot be copied*

The reason for this is that the copy-qubit operation is not linear and, hence, not unitary. Suppose we had a linear operation U_c that would copy a qubit. This means on state $(\alpha|0\rangle + \beta|1\rangle) \otimes |0\rangle$ it would do the following:

$$U_c[(\alpha|0\rangle + \beta|1\rangle) \otimes |0\rangle] = (\alpha|0\rangle + \beta|1\rangle) \otimes (\alpha|0\rangle + \beta|1\rangle) \tag{9}$$

$$= \alpha^2|00\rangle + \alpha\beta|01\rangle + \alpha\beta|10\rangle + \beta^2|11\rangle \tag{10}$$

On the other hand, since U_c is linear and because $(\alpha|0\rangle + \beta|1\rangle) \otimes |0\rangle = \alpha|00\rangle + \beta|10\rangle$:

$$U_c[\alpha|00\rangle + \beta|10\rangle] = \alpha|00\rangle + \beta|11\rangle \tag{11}$$

It is clear that (10) and (11) are the same if and only if $\alpha = 1$ and $\beta = 0$ or $\alpha = 0$ and $\beta = 1$. This is precisely the case if we have a classical 0 or 1. Hence, there cannot be a linear operation that copies an arbitrary unknown qubit.

Now imagine that Alice has an unknown qubit $|x\rangle = \alpha|0\rangle + \beta|1\rangle$ that she wants to send to Bob and that she furthermore can only communicate using classical bits. Is it in this case possible for Alice to communicate $|x\rangle$ to Bob? In the light of the no-cloning Theorem 1 it certainly is impossible to do this since whenever she measures x she will destroy/collapse it to a classical bit and she cannot copy it first. But suppose that Alice and Bob in addition each share one half of an EPR pair (8). The surprising observation is that there is a scheme that allows Alice to send or "teleport" $|x\rangle$ to Bob using only 2 classical bits [11].

In operational terms, the scheme works as follows. Let ϕ^+ be the first part of an EPR pair and ϕ^- the other half. That is, ϕ^+ is the first bit of $\frac{1}{\sqrt{2}}[|00\rangle + |11\rangle]$ and ϕ^- the second bit. Alice has ϕ^+ and Bob has ϕ^-. At some point Alice gets the unknown qubit $|x\rangle = \alpha|0\rangle + \beta|1\rangle$. She now does a unitary operation[1] on the two qubits, i.e., ϕ^+ and x. Then she measures these two qubits, obtaining two bits: 00, 01, 10, or 11. Next she send these two bits to Bob, who depending on the two bits, does one of four unitary operations on his ϕ^-. It turns out that this last unitary operation changes[2] ϕ^- into the unknown qubit $|x\rangle$. After the protocol, the EPR pair is destroyed, so in order to repeat this procedure a fresh EPR pair is needed.

[1] The unitary operation is a controlled-not of x on ϕ^+, followed by a Hadamard on x.

[2] In fact after the controlled-not and the Hadamard transform of Alice, it follows that their joint state is: $|00\rangle(\alpha|0\rangle+\beta|1\rangle)+|01\rangle(\alpha|1\rangle+\beta|0\rangle)+|10\rangle(\alpha|0\rangle-\beta|1\rangle)+|11\rangle(\alpha|1\rangle-\beta|0\rangle)$. This means that after Alice does her measurement, the third bit, i.e., ϕ^-, is the unknown qubit x up to a possible bit flip and/or phase shift depending on the outcome of Alice's measurement.

The important point for communication complexity is that this teleportation scheme is a way to simulate a qubit channel between Alice and Bob with a classical channel, at the cost of two bits per qubit, whenever Alice and Bob share EPR pairs.

Theorem 2. *[11] When Alice and Bob share EPR pairs, they can simulate a qubit channel with a classical bit channel at the cost of two classical bits per qubit.*

3 Quantum Black-Box Computation

Perhaps the simplest form of a computational task is the following. Suppose we have n Boolean variables $X_0, \ldots, X_{n-1}$, and we want to compute a property $P(X_0, \ldots, X_{n-1})$. The goal is to compute P looking at as few variables as possible. For example, suppose $P(X_0, \ldots, X_{n-1}) = 1$ iff there exists an i such that $X_i = 1$. That is, we want to compute the $OR(X_0, \ldots, X_{n-1})$. How many variables do we have to query? It is not too hard to see that we have to look at all the variables. A similar kind of reasoning shows that also in the randomized setting the bound is $\Omega(n)$. It has been shown by Grover [29] that a quantum algorithm can solve the OR with only $O(\sqrt{n})$ quantum queries.

Next we will turn our attention to another problem that allows even an exponential speedup. Define the following promise on the variables. We are guaranteed that they are either constant (i.e., all the X_i are either all 0 or all 1) or they are balanced: exactly half the X_i are 0 and the other half is 1. The problem is to find out whether the variables are constant or balanced.

It is easy to see that classically this problem requires $n/2 + 1$ queries to the variables. One of the first quantum algorithms, by Deutsch and Jozsa [33], establishes that this problem can be solved with just a single quantum query! Before we explain this algorithm we first have to explain how we model a quantum query.

Quantum Query. We have to model a quantum query in such a way that it is a unitary operation. We define a quantum query to variable X_i as follows. The state $|i, 0\rangle$ becomes after the query $|i, X_i\rangle$ and $|i, 1\rangle$ becomes $|i, 1 - X_i\rangle$. That is, for $1 \leq i \leq n$ and $b \in \{0, 1\}$:

$$|i, b\rangle \mapsto |i, b \oplus X_i\rangle$$

Since this describes what a query does on basis states, because of linearity it also works on states that are in superposition:

$$\sum_{i \in \{0,1\}^{\log(n)}, b \in \{0,1\}} \alpha_{i,b} |i, b\rangle \mapsto \sum_{i \in \{0,1\}^{\log(n)}, b \in \{0,1\}} \alpha_{i,b} |i, b \oplus X_i\rangle \ . \tag{12}$$

It can be easily checked that this operation is unitary.

The Deutsch-Jozsa Algorithm. Suppose n is a power of 2 and $l = \log n$. We start in a state with l 0s followed by a 1:

$$|0^l 1\rangle$$

Remember the Hadamard transform H on one qubit from (6). We do a Hadamard transform on all the qubits of the state, i.e., the following operation

$$\overbrace{H \otimes H \otimes \ldots \otimes H}^{l+1} = H^{\otimes l+1} \ .$$

This will result in the following state:

$$\frac{1}{\sqrt{n}} \sum_{i \in \{0,1\}^l} |i\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \tag{13}$$

Then we perform the only quantum query. This will affect our state according to (12) as follows:

$$\frac{1}{\sqrt{n}} \sum_{i \in \{0,1\}^l} (-1)^{X_i} |i\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \tag{14}$$

To see that this is correct, first observe that we perform the quantum query with the target qubit in superposition $(|0\rangle - |1\rangle)$ This means that $|i\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ after the query becomes $|i\rangle \frac{1}{\sqrt{2}}(|0 \oplus X_i\rangle - |1 \oplus X_i\rangle)$. Furthermore, if X_i is 0 then this is simply $|i\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$; on the other hand if $X_i = 1$ then it becomes $|i\rangle \frac{1}{\sqrt{2}}(|1\rangle - |0\rangle)$, which is the same as $(-1)|i\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. Hence, we get a factor of -1 iff $X_i = 1$. Next we apply again $H^{\otimes l+1}$ to the state and obtain the following messy-looking expression:

$$\frac{1}{\sqrt{n}} \sum_{i \in \{0,1\}^l} \frac{1}{\sqrt{n}} \sum_{j \in \{0,1\}^l} (-1)^{X_i \oplus (i,j)} |j\rangle |1\rangle \ , \tag{15}$$

where (i, j) is the inner-product between i and j modulo 2. Let us take a closer look at the part of this sum where $j = 0^l$:

$$\frac{1}{n} \sum_{i \in \{0,1\}^l} (-1)^{X_i} |0^l\rangle |1\rangle \tag{16}$$

Suppose that all the $X_i = 0$ and we are in the case "variables constant 0." Then (16) boils down to: $\frac{1}{n} \sum_{i \in \{0,1\}^l} |0^l\rangle |1\rangle = |0^l 1\rangle$. For the "constant 1" case we will end up in $(-1)|0^l 1\rangle$. This means that when we observe the final state in (15), we will see $0^l 1$ with probability 1. On the other hand, if half of the $X_i = 1$ and the other half are 0, then half of the terms in (16) are 1 and the other half are -1 and cancel each other out. The result of this is that $|0^l 1\rangle$ has amplitude 0 and will be seen with probability 0. So by observing state (15) we can conclude that if we observe $0^l 1$ we are in the constant case and if we observe anything else we are in the balanced case.

4 Applications in Distributed Computing

4.1 Communication

One of the main themes in quantum information processing is to extend classical communication and communication schemes with quantum ones. Here we will consider three models of quantum communication and compare them with classical communication.

1. Communication is done with qubits.
2. Both parties share EPR pairs but communication is done via a classical-bit channel.
3. Both parties share EPR pairs and communication is done with qubits.

The most simple form of communication is where Alice wants to send a message m of say k bits to Bob. We know that classically in general Alice needs to send k bits to Bob. Is this still true in the setting 1, 2, and 3? It follows from a theorem of Holevo [30] that when only qubits are used for communication Alice still needs to send k qubits. Moreover Cleve et al. [23] show that the same is true when both parties share EPR-pairs and classical communication is used.

For the third variant, where both EPR pairs and qubits are used, things are slightly different. Bennett and Wiesner [12] show that in this case there is a kind of a reverse of Theorem 2. This is a scheme, called super-dense coding, that allows Alice to send *two* classical bits with one qubit to Bob provided they share an EPR pair. It can be shown that, like Holevo's theorem, this is optimal.

4.2 Communication Complexity

Communication Complexity was introduced by Yao and Abelson [2,53]. Alice has an n-bit string x and Bob has an n-bit string y and their goal is to compute some function $f : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}$, minimizing the number of bits they communicate to each other. The area of communication complexity is well studied, see for example the books by Kushilevitz and Nisan [37] and Hromkovič [32]. The question we want to address here is: how does the communication complexity of certain problems vary when different models of quantum communication are used. We will denote $C(f)$ to denote the classical communication complexity of f. That is the number of bits the optimal protocol uses on the worst-case input. The model where only qubits can be used for communication (model 1, Sect. 4.1) was introduced by Yao [54]. We will use $Q(f)$ for the quantum communication complexity in the model where only qubits are used for communication. The first results in that model were lower bounds or impossibility results due to Yao and Kremer [36] and we will discuss them in Sect. 3.

The model where the communication is classical but both parties share entanglement, model 2, was introduced by Cleve and Buhrman [22]. We will denote the communication complexity in this model with $C^*(f)$, the model which uses both EPR pairs and qubits will be $Q^*(f)$. Cleve and Buhrman were the first to show that communication complexity can be reduced contrary to what one might

believe considering Holevo's theorem. Their setting differed slightly from the models we discuss here. In this setting they exhibit an example of a *three party* communication problem where the three parties share an entangled state, like an EPR pair but then for three parties. It is shown that when the parties share this entangled state the communication problem can be solved with two bits of communication whereas without such a prior shared state three bits are necessary. That is, there is a function f such that $C^*(f) = 2$ whereas $C(f) \geq 3$. Better separations in the multiparty setting were found in [15] and [21]. The latter paper exhibits a function f for k parties such that $C^*(f) = k$ and $C(f) = \Omega(k \log(k))$.

Next we will turn our attention to the qubit communication model $Q(f)$. However, keep in mind that protocols for this model can be translated to the model where both parties share EPR pairs and communicate classically, since via teleportation, Theorem 2 gives us: $C^*(f) \leq 2Q(f)$.

Deutsch-Jozsa Communication Problem. The first gap for two-party qubit communication complexity was demonstrated by Buhrman, Cleve, and Wigderson [17]. They showed for a promise version of the equality problem[3], EQ', that $Q(EQ') = O(\log(n))$ and that also $C(EQ') = \Omega(n)$. This exhibits an exponential gap between classical and quantum communication complexity. The quantum protocol is inspired by the Deutsch-Jozsa algorithm from Sect. 3 and the classical bound stems from a deep and surprising combinatorial theorem from Frankl and Rödl [27].

$EQ'(x, y) = 1$ iff $x = y$ but with the extra promise that it will always be the case that the Hamming distance $\Delta(x, y) = 0$ or $n/2$. The Hamming distance between two strings x and y, $\Delta(x, y)$, is the total number of bits where x and y are different. We will see that EQ' can be solved with just $\log(n) + 1$ qubits of communication from Bob to Alice. Note that under the Hamming distance promise, Alice and Bob have to figure out whether $x_1 \oplus y_1 \ldots x_n \oplus y_n$ is constant or balanced, since in the constant 0 case $x = y$ and in the balanced $x \neq y$. So if we set $X_i = x_i \oplus y_i$ then we have the Deutsch-Jozsa problem back.

If Alice could obtain the final state from equation (15),

$$\frac{1}{\sqrt{n}} \sum_{i \in \{0,1\}^l} \frac{1}{\sqrt{n}} \sum_{j \in \{0,1\}^l} (-1)^{X_i \oplus (i,j)} |j\rangle |1\rangle \ ,$$

she would do a final measurement and know the answer. To this end Bob prepares the following state:

$$\frac{1}{\sqrt{n}} \sum_{i \in \{0,1\}^l} |i\rangle \frac{1}{\sqrt{2}} (|0 \oplus y_i\rangle - |1 \oplus y_i\rangle)$$

[3] $EQ(x, y) = 1$ if $x = y$ and 0 otherwise. EQ requires n bits of communication. A promise version of a problem means that Alice and Bob are only required to compute the answer correctly on certain instances that fall within the promise and it doesn't matter what they compute on the other instances that don't satisfy the promise.

and sends these $\log(n)+1$ qubits to Alice. Alice then performs the unitary transformation that changes state $|i\rangle|b\rangle$ to $|i\rangle|b \oplus x_i\rangle$ resulting in state:

$$\frac{1}{\sqrt{n}} \sum_{i \in \{0,1\}^l} |i\rangle \frac{1}{\sqrt{2}}(|0 \oplus y_i \oplus x_i\rangle - |1 \oplus y_i \oplus x_i\rangle)$$

which is after we rewrite it *precisely* the state from (14):

$$\frac{1}{\sqrt{n}} \sum_{i \in \{0,1\}^l} (-1)^{X_i} |i\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$$

Next Alice proceeds as in the Deutsch-Josza algorithm and applies $H^{\otimes \log(n)+1}$ and measures the final state.

The general idea is to use a quantum black-box algorithm in a distributed setting. Whenever the black-box algorithm wants to make a query, Alice and Bob exchange a round of $\log(n)+1$ qubits and Alice continues the black-box algorithm. This allows one in general to use any black-box algorithm as a communication protocol. In this way it can be shown that, by using Grover's algorithm [29] the Disjointness problem can be solved with $O(\sqrt{n}\log(n))$ many qubits [17].

Bounded-Error Protocols. All the above (quantum) protocols don't make errors and compute the outcome exactly. When studying randomized versions of communication complexity, however, it is unavoidable to introduce errors. A classical randomized protocol for f, $R_2(f)$, is a protocol where both Alice and Bob can use random bits. They are required to compute the correct outcome with probability at least $2/3$. The distinction between private and public random bits can be made, where in the public bit/coin model Alice and Bob see the same random bits and in the private they each have a different random source. Newman [45] has shown that up to an additive logarithmic term the models are the same.

Rabin and Yao show for EQ that there exists a classical randomized protocol that only needs $O(\log(n))$ bits: $R_2(EQ) = O(\log(n))$. This implies that the promise problem EQ' also has a $O(\log(n))$ randomized classical bit protocol that is correct with probability at least $2/3$. Note, however, that the quantum protocol never makes an error.

The disjointness problem $DISJ$ is defined as follows. Alice and Bob each have a subset A and B of $\{0,1\}^n$, they have to decide whether $A \cap B = \emptyset$. Kalyanasundaram and Schnitger [34] show that this problem also has high communication complexity in the randomized setting: $R_2(DISJ) = \Omega(n)$. Buhrman et al. in the same paper show that when we allow the quantum protocol to compute the answer with probability at least $2/3$, we denote this by $Q_2(f)$, that $Q_2(DISJ) = O(\sqrt{n}\log(n))$. This bound was improved by [31] and a $O(\sqrt{n})$ protocol was recently constructed by Aaronson and Ambainis [1]. Razborov has shown, using some variant of the polynomial method, that this bound is tight [49]

The disjointness problem demonstrates a quadratic gap between classical randomized and quantum communication complexity. Moreover this is an example of a gap known where the function f is not a promise problem. The only other known total problem that allows for a more efficient quantum protocol is that of the equality problem in the *simultaneous message passing* model [16]. The main ingredients to the protocol are a quantum fingerprinting scheme and a test to distinguish orthogonal states from parallel ones. In the simultaneous message passing model Alice and Bob don't send messages to each other but send one message to a third party, called the referee. The referee only sees the messages from Alice and Bob and has to output $f(x, y)$.

The biggest gap between the randomized and the quantum two party communication complexity model was obtained by Raz [48]. He showed that there is a promise problem f such that $Q(f) = O(\log(n))$ but $R_2(f) = \Omega(\sqrt{n})$. Ambainis et al. [7] also exhibit an exponential gap between quantum protocols and classical protocols for a different form of communication problem called sampling which we shall not discuss here further.

Summarizing for promise problems there exist exponential gaps between classical and quantum communication complexity. For total problems the best known gap is only quadratic. In turn this sheds some light on the EPR paradox. Holevo's theorem proves that EPR pairs cannot be used to reduce communication. Since all the protocols in this section work for the model where the parties share EPR pairs and communicate classically it follows that EPR pairs can reduce the communication complexity of certain problems. This situation seems contradictory but notice that the actual amount of information that needs to be communicated between Alice and Bob is only 1 bit, namely the outcome of f.

Lower Bounds. In the previous section we showed that quantum communication protocols are sometimes superior to classical protocols. In this section we examine the converse and turn our attention to lower bounds for quantum communication complexity.

Classically for deterministic communication complexity there is a general technique for proving lower bounds. For any function $f : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}$ one can define the boolean $2^n \times 2^n$ communication matrix $M_f(x, y) = f(x, y)$. Mehlhorn and Schmidt [44] related the rank of this matrix to the communication complexity. They show that $\log(\text{rank}(M_f)) \leq C(f)$. This is a very useful tool. Take for example the equality problem. The communication complexity matrix for EQ is the $2^n \times 2^n$ identity matrix which has only 1's on the diagonal and is 0 on off-diagonal entries. Since this matrix has rank 2^n it follows that $C(f) \geq n$. A similar statement is true in the quantum setting:

Theorem 3. *For any communication problem f:*

1. $\log(\textit{rank}(M_f))/2 \leq Q(f)$ *[36].*
2. $\log(\textit{rank}(M_f)) \leq C^*(f)$ *[18].*
3. $\log(\textit{rank}(M_f))/2 \leq Q^*(f)$ *[18].*

A natural and long standing open problem is whether the communication complexity is also a lower bound for the log-rank. That is, whether the log-rank characterizes the communication complexity. The biggest known gap between the log-rank and the communication complexity is almost quadratic [47]. The log-rank conjecture states that for every total f, $\log(\text{rank}(f))$ and $C(f)$ are all polynomially related. It follows from Theorem 3 that if the log-rank conjecture is true then for total f: $Q(f), C^*(f), Q^*(f)$, and $C(f)$ are polynomially related.

The log rank lower bound method only works well for errorless protocols. For bounded error models there is another bound called *discrepancy*. Kremer [36] and Yao show that the discrepancy bound also works for the bounded error qubit communication model Q_2. This enables them to show a linear lower bound in this model for a problem called inner product modulo 2, IP. Here $IP(x, y) = x_1 \cdot y_1 + \cdots + x_n \cdot y_n \mod 2$. Ambainis et al. [7] extend this bound to also yield a $\Omega(n)$ bound even when Alice and Bob are allowed to make an error which is very close to $1/2$.

For the model where both parties share EPR pairs, Cleve et al. [23] were the first to show a linear lower bound for IP. They came up with a new technique that is essentially quantum mechanical in nature. It can be seen as a quantum adversary argument. This enabled them to show that any (quantum) protocol for IP can be (ab)used, when run in superposition, to communicate n bits from Alice to Bob. Let $Q_2^*(f)$ denote the communication complexity of f where Alice and Bob compute f correctly with probability $2/3$, they share EPR pairs and the communication is with qubits.

Theorem 3 yields a lower bound of $\Omega(n)$ for $DISJ$ in the errorless models since the M_{DISJ} has rank 2^n. In the bounded error model recently Razborov showed, in a very nice paper, that the $DISJ$ needs $\Omega(\sqrt{n})$ qubits of communication even in the presence of shared EPR pairs. Summarizing we have the following theorem:

Theorem 4. *1.* $Q_2^*(IP) = \Omega(n)$ *[36,23].*
2. $Q_2^*(DISJ) = \Omega(\sqrt{n})$ *[49]*

4.3 Loopholes in Nonlocality Experiments

Tools and results from the study of quantum communication complexity have been applied fruitfully to tune parameters in physical experiments that test the "quantumness" of our world. The EPR paradox has been and still is a subject of dispute. Much progress was made when Bell [10] came up with a test that would, in case quantum mechanics was correct, show correlations that could not be explained with just classical reasoning. Such nonlocality experiments have been performed in the lab and non-classical correlations have been observed [9].

However, experimental realizations of the nonlocality tests are hampered by noise and imperfections in the physical apparatus. In particular, measurement devices for individual quantum systems (e.g., single-photon detectors) tend to fail on most runs of the experiment, allowing local classical explanations of the data by means of local classical theories that are allowed to make the same kind

of errors and this opens the so-called "detection loophole." Ideas from quantum communication complexity have been used by Brassard et al. [14], Massar [40], and Buhrman et al. [19] to propose new nonlocality experiments and to bound the maximum detector efficiency, minimum noise, and hidden communication using which the results can be explained by means of a classical local model. The goal is to construct an experiment that demonstrates the nonlocal character of quantum mechanics even when the experiments are faulty and make errors.

An experiment is modeled as two (or more) parties Alice and Bob that each have an input of length n. However, contrary to the communication complexity model, Alice and Bob are not allowed to communicate with each other. In the classical setting Alice and Bob share a common source of random bits, and in the quantum scenario Alice and Bob share EPR pairs or more generally an entangled state. Alice now will depending on her input and her random bits (or some operation on her part of the EPR pairs) output some string a of m bits. Bob follows some protocol to also output m bits b. This way they produce correlation distributions $\Pr[a, b \mid x, y]$. The goal now is to come up with a set of correlation distributions and show that there is a quantum protocol that generates these distributions whereas every classical protocol fails to do so even if it is allowed to make small errors or sometimes not produce an output at all.

Deutsch-Josza Correlations. To demonstrate these ideas, we return once more to the Deutsch-Josza problem, following Brassard et al. [14] and Massar [40]. This time, Alice and Bob cannot communicate, but they start out sharing a quantum state, receive classical bit strings $x, y \in \{0,1\}^n$, respectively; both Alice and Bob produce outputs, $a, b \in \{0,1\}^l$, respectively, and we are interested in the *correlations* between these outputs, namely the probability distributions $\Pr[a, b \mid x, y]$ of Alice outputting a and Bob outputting b given that Alice got input x and Bob input y. Recall that the "trick" in turning the Deutsch-Josza algorithm into a communication protocol was to let Bob perform the first steps of the algorithm and then send the quantum state to Alice who completed the steps with her input. Now, since Alice and Bob cannot communicate, we replace the quantum channel by EPR pairs. Alice and Bob start out with the following state comprised of $l = \log(n)$ EPR pairs and two auxiliary qubits:

$$\frac{1}{2\sqrt{n}} \sum_{i \in \{0,1\}^l} |i\rangle \left(|0\rangle - |1\rangle\right) |i\rangle \left(|0\rangle - |1\rangle\right)$$

Here, Alice has the first $l+1$ qubits and Bob the remaining $l+1$ qubits. Now they pretend that each on her/his side are in the Deutsch-Jozsa algorithm before the oracle query, as given in (13). Accordingly, they perform the operation $|i\rangle|b\rangle \mapsto |i\rangle|b \oplus y_i\rangle$ on their part of the state, resulting in the following global state:

$$\frac{1}{2\sqrt{n}} \sum_{i \in \{0,1\}^l} (-1)^{x_i + y_i} |i\rangle \left(|0\rangle - |1\rangle\right) |i\rangle \left(|0\rangle - |1\rangle\right)$$

Then they apply the Hadamard operation on their $l+1$ qubits, yielding the state

$$\frac{1}{n\sqrt{n}} \sum_{i\in\{0,1\}^l} (-1)^{x_i+y_i} \left(\sum_{a\in\{0,1\}^l} (-1)^{(i,a)}|a\rangle \right) |1\rangle \left(\sum_{b\in\{0,1\}^l} (-1)^{(i,b)}|b\rangle \right) |1\rangle$$

$$= \frac{1}{n\sqrt{n}} \sum_{a,b\in\{0,1\}^l} \left(\sum_{i\in\{0,1\}^l} (-1)^{x_i+y_i+(i,a\oplus b)} \right) |a\rangle|1\rangle|b\rangle|1\rangle$$

Now they both measure and output their measurement. By the laws of quantum mechanics, the probability for Alice to observe $|a\rangle|1\rangle$ and Bob $|b\rangle|1\rangle$ is

$$\Pr[a, b \mid x, y] = \frac{1}{n^3} \left(\sum_{i\in\{0,1\}^l} (-1)^{x_i+y_i+(i,a\oplus b)} \right)^2$$

If $x = y$, then

$$\Pr[a, b \mid x, y] = \begin{cases} \frac{1}{n} & \text{if } a = b \\ 0 & \text{if } a \neq b \end{cases}$$

whereas for $\Delta(x, y) = n/2$ and $a = b$ we have $\Pr[a, b \mid x, y] = 0$. Hence, the outputs are correlated in that whenever $x = y$, we always see $a = b$ and whenever $\Delta(x, y) = n/2$, we never see $a = b$.

Can these correlations be realized by a classical protocol with shared randomness and no communication? No, since then Bob could send his output to Alice, solving the communication problem with $O(\log n)$ bits, which is ruled out by the lower bound of $\Omega(n)$. Then, how closely can they be realized approximately, i.e., how precise does an experiment need to be? For the "detection loophole," it is assumed that any measurement succeeds with probability at least η and if it fails, there will be no output. Then η^2 is the probability that both Alice's and Bob's measurements succeed. If the world is classical, then we have an adversary who is trying to reproduce the correlations without communication using the possibility not to produce an output on a η^2 fraction of the runs of the experiment. By the Yao principle there will be for any distribution on the inputs a classical local *deterministic* strategy which produces a (correct) output for an η^2 fraction of the inputs. Consider the input distribution where $x \in \{0, 1\}^n$ is chosen uniformly and random and $y = x$; fix the best deterministic strategy. Let $Z_a = \{x : \text{Alice and Bob output } a\}$, then

$$\eta^2 2^n \leq \sum_{a\in\{0,1\}^l} |Z_a|$$

Moreover, for each $a \in \{0, 1\}^l$, $Z_a \subseteq \{0, 1\}^n$ must not contain x, y with $\Delta(x, y) = n/2$, therefore, by a deep theorem by Frankl and Rödl [27], $|Z_a| \leq 2^{0.993n}$. This implies $\eta^2 2^n \leq n2^{0.993n}$ or $\eta \leq \sqrt{n2^{-0.007n}}$. Hence, with growing n, the detector efficiency at which there still exists a classical local model decreases exponentially. So if the quality of the measurement equipment does not decrease too fast with growing n, the detection loophole can be "closed" with an experiment for the Deutsch-Jozsa correlations.

There are several issues with this approach. In a nonlocality experiment, the input distribution should a product distribution so that it can be implemented locally in the lab. Furthermore, there are very efficient classical *bounded-error* protocols for equality, implying that the quantum correlations above can be very well simulated classically if the experiment is subject to noise. And finally, an asymptotic analysis is often too coarse since the region where the bounds kick in may be out of reach experimentally.

Concerning the bounded-error case, a multiparty nonlocality experiment has been constructed, building again on an earlier multiparty quantum communication protocol [21]. This family of experiments has $\eta \leq 1/k^{1/6}$ and tolerates error $1/2 - 1/\,\mathrm{o}(k^{1/6})$, where k is the number of parties [20].

4.4 Coin Tossing

Research into quantum cryptography is motivated by two observations about quantum mechanics:

1. Nonorthogonal quantum states cannot be distinguished perfectly and parts of certain orthogonal quantum states cannot be distinguished if the remaining parts are inaccessible;
2. Measurement disturbs the quantum state. This is the so-called "collapse of the wave function."

The second observation hints at the possibility of detecting eavesdroppers or other types of cheaters, whereas the first property appears to allow hiding data, both unhampered by unproven computational assumptions. Indeed, for the task of cooperatively establishing a random bit string between two parties in the presence of eavesdroppers, quantum key distribution [13,41,39] achieves security against the most general attack by an adversary that has unbounded computational power but has to obey the laws of quantum mechanics.

Initially, it was thought that these properties would admit protocols for the cryptographic primitive "bit commitment." In bit commitment, there are two parties Alice and Bob; in the initial phase of the protocol, Alice has a bit b and communicates with Bob to "commit" to the value of b without revealing it. At a later time, Alice "unveils" her bit, allowing Bob to perform checks against the information obtained in the initial phase. The properties sought of bit-commitment protocols are that they are "concealing" (Bob does not learn anything about b in the initial phase) and "binding" (Bob will catch Alice trying to unveil $1 - b$ instead of b).

Unfortunately, Mayers [42] and Lo and Chau [38] proved that perfect quantum bit commitment is impossible. Their impossibility result extends to "coin tossing" [43,38], a weaker cryptographic primitive where the two parties want to agree on a random bit whose value cannot be influenced by either of them. Moreover, the impossibility extends even to the case of "weak coin tossing" [4], where outcome $b = 0$ is favorable for Alice and outcome $b = 1$ favorable for Bob, thus ruling out perfect quantum protocols for leader election. However, what

turned out to be possible are coin-tossing protocols, where there are guarantees on how much a cheater can bias the outcome.

Consider k parties out of which at most $k' < k$ are dishonest; which players are dishonest is fixed in advance but unknown to the honest players. The players can communicate over broadcast channels. Initially they do not share randomness, but they can privately flip coins; the probabilities below are with respect to the private random coins. A coin-flipping protocol establishes among the honest players a bit b such that

- if all players are honest, $\Pr[b = 0] = \Pr[b = 1] = 1/2$
- if up to k' players are dishonest, then $\Pr[b = 0], \Pr[b = 1] \leq 1/2 + \epsilon$

ϵ is called the *bias*; a small bias implies that colluding dishonest players cannot strongly influence the outcome of the protocol. Players may abort the protocol.

Classically, if a (weak) majority of the players is bad then no bias $< 1/2$ can be achieved and hence no meaningful protocols exist [50]. For example, if we only have two players and one of them is dishonest, then no protocols with bias $< 1/2$ exist. (For a minority of bad players, quite nontrivial protocols exist; see [26].) Allowing quantum bits (qubits) to be sent instead of classical bits changes the situation dramatically. Surprisingly, in the two-party case coin flipping with bias $< 1/2$ is possible, as was first shown in [3]. The best known bias is $1/4$ and this is optimal for a special class of three-round protocols [4]; for a bias of ϵ at least $\Omega(\log\log(1/\epsilon))$ rounds of communication are necessary [4]. Recently, Kitaev (unpublished, see [35,6]) showed that in the two-party case no bias smaller than $1/\sqrt{2} - 1/2$ is possible.

In the weak version of the coin-flipping problem, we know in advance that outcome 0 benefits Alice and outcome 1 benefits Bob. In this case, we only need to bound the probabilities of a dishonest Alice convincing Bob that the outcome is 0 and a dishonest Bob convincing Alice that the outcome is 1. In the classical setting, a standard argument shows that even weak coin flipping with a bias $< 1/2$ is impossible when a majority of the players is dishonest. In the quantum setting, this scenario was first studied for two parties under the name *quantum gambling* [28]. Subsequently, Spekkens and Rudolph [51] gave a quantum protocol for two-party weak coin flipping with bias $1/\sqrt{2} - 1/2$ (i.e., no party can achieve the *desired outcome* with probability greater than $1/\sqrt{2}$). Notice that this is a better bias than in the best strong coin flipping protocol of [4]. Kitaev's lower bound for strong coin flipping does not apply to weak coin flipping. Thus, weak protocols with arbitrarily small $\epsilon > 0$ may be possible. The only known lower bounds for weak coin flipping are that the protocol of [51] is optimal for a restricted class of protocols [5] and that a protocol must use at least $\Omega(\log\log(1/\epsilon))$ rounds of communication to achieve bias ϵ (shown in [4] for strong coin flipping but the proof also applies to weak coin flipping).

Quantum coin flipping and leader election for more than two parties were investigated by Ambainis et al. [6]: Even if there is only a single honest party among k players, bias $1/2 - c/k^{1.78}$ can still be achieved by a quantum protocol (for some $c > 0$) and there is a lower bound that for some $c' > 0$, $1/2 - c'/k$ cannot be achieved. Both bounds can be generalized to the situation where at

most $(1-\epsilon)k$ of the players are bad, for $\epsilon > 0$; in this case, bias $\delta < 1/2 - c''\epsilon^{1.78}$ is achievable independent of the number of players and achieving constant bias $\delta < 1/2 - c'''\epsilon$ is impossible, for constants $c'', c''' > 0$.

5 Conclusion and Open Problems

We have surveyed some of the results in quantum distributed computing. Many problems however remain. What is the relationship between the various models, Q, C^*, Q^* both in the errorless and in the bounded error setting? For the errorless models, a positive answer to the log-rank conjecture shows that they are all polynomially related but also this is at the moment still wide open.

We have seen that exponential gaps between classical and quantum communication complexity problems are possible, however, all of these examples entailed promise problems. Can there also be exponential gaps for total problems in the bounded error setting?

Techniques and protocols from quantum and classical communication complexity can help to construct nonlocality experiments. It remains an open question what the best bounds for two and more parties are for the error of the experiment and the detector efficiency.

A question that is sheds some light on the relationship between Q, C^*, and Q^* is the following. Given a correlation game with two parties that each get inputs of size n and produce outputs of size m, and use an entangled state of a finite amount of qubits. Is there a protocol that uses only an $O\log(n+m))$ qubit entangled state that can be used to approximate, say in terms of small total variation distance, the correlations from the original protocol? Such a statement is true in the classical scenario with respect to the number of shared random bits and has a very similar prove of the fact that for any communication complexity protocol only $O(\log(n))$ shared random bits are needed [45]. Note that if it can be shown that $O(\log(n))$ entangled qubits are enough to simulate an arbitrary communication complexity protocol on inputs of length n then C_2^*, Q_2^*, and Q_2 are all related with an additional overhead of $O(\log(n))$ qubits of communication.

In the simultaneous message passing model there exists a $O(\log n)$ protocol that solves the equality problem. The equality problem is equivalent to the problem of deciding whether $x \oplus y = 0$, where $x \oplus y$ is the bitwise XOR of x and y. No such protocol is known for the three party problem to decide whether $x \oplus y \oplus z = 0$. The best known quantum protocol is due to Ambainis and Shi [8] who need $O(\sqrt{n})$ qubits to solve this problem. However, the best known lower bound for this three party problem in the classical setting is $\Omega(\sqrt{n})$ and the best known classical upper bound is $O(n^{2/3})$

Quantum bit commitment and perfect coin tossing have been shown to be impossible, but there are protocols for coin tossing with constant bias. The best known impossibility bound for strong coin tossing matches the bias of the best known protocol for weak coin tossing – it is not clear whether this is a coincidence. Tight bounds on the achievable bias are not known in both cases; we even do not know whether there exists a protocol with a finite number of rounds and qubits that guarantees the optimal bias, or whether there are more and more complex protocols whose biases converge.

References

1. S. Aaronson and A. Ambainis. Quantum search of spatial regions. quant-ph/0303041, 2003.
2. H. Abelson. Lower bounds on information transfer in distributed computations. *J. Assoc. Comput. Mach.*, 27(2):384–392, 1980. Earlier version in FOCS'78.
3. D. Aharonov, A. Ta-Shma, U. Vazirani, and A. Yao. Quantum bit escrow. In *Proceedings of STOC'00*, pages 705–714, 2000.
4. A. Ambainis. A new protocol and lower bounds for quantum coin flipping. In *Proceedings of 33rd ACM STOC*, pages 134–142, 2001.
5. A. Ambainis. Lower bound for a class of weak quantum coin flipping protocols. quant-ph/0204063, 2002.
6. A. Ambainis, H. Buhrman, Y. Dodis, and H. Röhrig. Multiparty quantum coin flipping. Submitted, 2003.
7. A. Ambainis, L. Schulman, A. Ta-Shma, U. Vazirani, and A. Wigderson. The quantum communication complexity of sampling. In *39th IEEE Symposium on Foundations of Computer Science*, pages 342–351, 1998.
8. A. Ambainis and Y. Shi. Distributed construction of quantum fingerprints. quant-ph/0305022, 2003.
9. A. Aspect, J. Dalibard, and G. Roger. Experimental test of Bell's inequalities using time-varying analyzers. *Phys. Rev. Lett.*, 49(25):1804, 1982.
10. J. S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1, 1964.
11. C. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physiscal Review Letters*, 70:1895–1899, 1993.
12. C. Bennett and S. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physiscal Review Letters*, 69:2881–2884, 1992.
13. C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984.
14. G. Brassard, R. Cleve, and A. Tapp. The cost of exactly simulating quantum entanglement with classical communication. *Physical Review Letters*, 83(9):1874–1877, 1999.
15. H. Buhrman, R. Cleve, and W. van Dam. Quantum entanglement and communication complexity. *SIAM Journal on Computing*, 30(8):1829–1841, 2001. quant-ph/9705033.
16. H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16), September 26, 2001.
17. H. Buhrman, R. Cleve, and A. Wigderson. Quantum vs. classical communication and computation. In *30th Annual ACM Symposium on Theory of Computing*, 1998. quant-ph/9702040.
18. H. Buhrman and R. de Wolf. Communication complexity lower bounds by polynomials. In *16th IEEE Annual Conference on Computational Complexity (CCC'01)*, pages 120–130, 2001. cs.CC/9910010.
19. H. Buhrman, P. Høyer, S. Massar, and H. Röhrig. Combinatorics and quantum nonlocality. Accepted for publication in Physical Review Letters, 2002.
20. H. Buhrman, P. Høyer, S. Massar, and H. Röhrig. Resistance of quantum nonlocality to imperfections. Manuscript, 2003.
21. Harry Buhrman, Wim van Dam, Peter Høyer, and Alain Tapp. Multiparty quantum communication complexity. *Physical Review A*, 60(4):2737 – 2741, October 1999.

22. R. Cleve and H. Buhrman. Substituting quantum entanglement for communication complexity. *Physical Review A*, 56(2):1201–1204, august 1997.
23. R. Cleve, W. van Dam, M. Nielsen, and A. Tapp. Quantum entanglement and the communication complexity of the inner product function. In Springer-Verlag, editor, *Proceedings of the 1st NASA International Conference on Quantum Computing and Quantum Communications*, 1998.
24. D. Dieks. Communication by EPR devices. *Phys. Lett. A*, 92(6):271–272, 1982.
25. A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777, 1935.
26. U. Feige. Noncryptographic selection protocols. In *Proceedings of 40th IEEE FOCS*, pages 142–152, 1999.
27. P. Frankl and V. Rödl. Forbidden intersections. *Trans. Amer. Math. Soc.*, 300(1):259–286, 1987.
28. L. Goldenberg, L. Vaidman, and S. Wiesner. Quantum gambling. *Physical Review Letters*, 88:3356–3359, 1999.
29. L. Grover. A fast quantum mechenical algorithm for database search. In *28th ACM Symposium on Theory of Computing*, pages 212–218, 1996.
30. A. S. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii*, 9(3):3–11, 1973. English translation in *Problems of Information Transmission*, 9:177–183, 1973.
31. P. Høyer and R. de Wolf. Improved quantum communication complexity bounds for disjointness and equality. In *Proceedings of 19th Annual Symposium on Theoretical Aspects of Computer Science (STACS'2002)*, volume 2285 of *Lecture Notes in Computer Science*, pages 299–310. Springer, 2002. quant-ph/0109068.
32. J. Hromkovič. *Communication Complexity and Parallel Computing*. EATCS series: Texts in Theoretical Computer Science. Springer, 1997.
33. D. Deutsch R. Josza. Rapid solutions of problems by quantum computation. *Proc. Roy. Soc. London Se. A*, 439:553–558, 1992.
34. B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM J. Discrete Mathematics*, 5(4):545–557, 1992.
35. A. Yu. Kitaev. Quantum coin-flipping. Talk at QIP 2003 (slides and video at MSRI), December 2002.
36. I. Kremer. Quantum communication. Master's thesis, Computer Science Department, The Hebrew University, 1995.
37. E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
38. H. K. Lo and H. F. Chau. Why quantum bit commitment and ideal quantum coin tossing are impossible. *Physica D*, 120:177–187, 1998.
39. H-K. Lo and H. F. Chau. Unconditional security of quantum key distribution over arbitrarily long distances. quant-ph/9803006, 3 Mar 1998.
40. S. Massar. Nonlocality, closing the detection loophole, and communication complexity. *Physical Review A*, 65:032121, 2002.
41. D. Mayers. Unconditional security in quantum cryptography. quant-ph/9802025, 10 Feb 1998.
42. D. Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters*, 78:3414–3417, 1997.
43. D. Mayers, L. Salvail, and Y. Chiba-Kohno. Unconditionally secure quantum coin tossing. quant-ph/9904078, 22 Apr 1999.

44. K. Mehlhorn and E. M. Schmidt. Las Vegas is better than determinism in VLSI and distributed computing (extended abstract). In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*, pages 330–337, San Francisco, California, 5–7 May 1982.
45. I. Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, 39(2):67–71, July 1991.
46. M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
47. N. Nisan and A. Wigderson. On rank vs. communication complexity. *Combinatorica*, 15:557–566, 1995. Earlier version in FOCS'94.
48. R. Raz. Exponential separation of quantum and classical communication complexity. In *Proceedings of 31th STOC*, pages 358–367, 1999.
49. A. A. Razborov. Quantum communication complexity of symmetric predicates. *Izv. Math.*, 67(1):145–159, 2003.
50. M. Saks. A robust noncryptographic protocol for collective coin flipping. *SIAM J. Discrete Math.*, 2(2):240–244, 1989.
51. R. Spekkens and T. Rudolph. A quantum protocol for cheat-sensitive weak coin flipping. quant-ph/0202118, 2002.
52. W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.
53. A. C-C. Yao. Some complexity questions related to distributive computing. In *Proceedings of 11th STOC*, pages 209–213, 1979.
54. A. C-C. Yao. Quantum circuit complexity. In *Proceedings of 34th FOCS*, pages 352–360, 1993.